

WHAT IS CLAIMED IS:

1. A computer-readable medium having computer-executable instructions, comprising:

receiving a request to open a file, the file having state
5 information associated therewith;

in response to the request, sending a first handle to use
to access data in the file and at least part of the state
information associated therewith, the first handle having
access rights to the file;

10 sending a resume key by which a duplicate handle may be
requested, the duplicate handle having access rights to the
file that correspond to the first handle;

receiving the resume key and a request for a duplicate
handle to the file;

15 sending the duplicate handle in response to the request
for a duplicate handle; and

providing access to the file via the duplicate handle.

2. The computer-readable medium of claim 1, wherein the
20 state information includes a mode in which the file is opened.

3. The computer-readable medium of claim 2, wherein the mode includes at least one of read only, read/write, a lock, and a mode in which the file is opened for exclusive use by the handle and any duplicates thereof.

5

4. The computer-readable medium of claim 3, wherein the lock comprises a range that indicates bytes of the file that can only be accessed by the first handle and any duplicates thereof.

10

5. The computer-readable medium of claim 1, further comprising authenticating a client that sends any request.

6. The computer-readable medium of claim 5, further
15 comprising authenticating a server that sends any response.

7. The computer-readable medium of claim 1, wherein the request to open a file is sent over a first channel and the request for a duplicate handle to the file is sent over a
20 second channel.

8. The computer-readable medium of claim 7, wherein unencrypted information is sent over the first channel and encrypted information is sent over the second channel.

5 9. The computer-readable medium of claim 7, wherein the channels pass through a single network interface on a single client.

10 10. The computer-readable medium of claim 7, wherein the channels pass through at least two interfaces on a single client.

15 11. The computer-readable medium of claim 7, wherein the first channel is disconnected and wherein the file remains open afterwards.

20 12. The computer-readable medium of claim 11, wherein the resume key is received after the first channel has become disconnected.

13. The computer-readable medium of claim 7, wherein the first channel is disconnected and wherein the file remains

open afterwards at least until another client requests access to the file.

14. The computer-readable medium of claim 13, wherein if
5 the other client requests access to the file before a time has expired, the other client is denied access to the file.

15. The computer-readable medium of claim 13, wherein if
the other client requests access to the file after the time
10 has expired, the file is closed and the other client is granted access to the file.

16. The computer-readable medium of claim 15, wherein
the client is informed that the file has been accessed by
15 another client after the client sends the resume key with a request for a duplicate handle.

17. The computer-readable medium of claim 13, wherein
after the resume key is received, the client is informed that
20 the file has not been changed by another client.

18. The computer-readable medium of claim 1, wherein the request to open the file is sent from a first client and the duplicate handle is sent to a second client.

5 19. The computer-readable medium of claim 18, wherein the first client sends the duplicate handle to the second client.

20. The computer-readable medium of claim 18, wherein
10 the second client receives the duplicate handle from a server that provides access to the file.

21. A computer-readable medium having computer-executable instructions, comprising:

15 sending a request to open a file in a file access mode including data indicative of access rights;

 receiving a first handle to use to access data in the file in accordance with the access rights, the first handle providing access to all or part of any state information
20 associated with the file;

sending a resume key with a request for a duplicate
handle, the duplicate handle having access rights that
correspond to the access rights of the first handle;
receiving the duplicate handle; and
5 accessing the file via the duplicate handle.

22. The computer-readable medium of claim 21, further
comprising requesting the resume key.

10 23. The computer-readable medium of claim 21, wherein
the resume key is automatically returned.

24. The computer-readable medium of claim 21, wherein
the file is accessed via the first handle over a first channel
15 and the file is accessed via the duplicate handle over a
second channel.

25. The computer-readable medium of claim 24, wherein
the second channel comprises a remote direct memory access
20 (RDMA) channel in which data can be transferred from a client
to a server without assistance from a central processing unit
(CPU) on either the client or the server.

26. The computer-readable medium of claim 24, wherein the first channel is closed before the second channel is established.

5

27. The computer-readable medium of claim 26, wherein the file remains open after the first channel is closed.

28. The computer-readable medium of claim 21, wherein
10 the file access mode comprises a mode of exclusive use by the handle and any duplicates thereof.

29. The computer-readable medium of claim 21, wherein the mode comprises an exclusive lock on a range of bytes in
15 the file.

30. The computer-readable medium of claim 29, wherein the range of bytes is accessed via the duplicate handle.

20 31. A computer-readable medium having computer-executable instructions, comprising:

opening a file in a mode comprising a set of one or more
access rights and associating state information therewith;

obtaining a first handle to the file, the first handle
capable of accessing the file in accordance with any access
5 right included in the mode, the first handle obtained via a
first channel;

obtaining a resume key with which to request a duplicate
handle, the duplicate handle capable of accessing the file in
any way in which the first handle is capable;

10 closing the first channel and keeping the file opened
afterwards;

sending a request for the duplicate handle together with
the resume key, the request being sent via a second channel;
and

15 accessing the file via the duplicate handle.

32. The computer-readable medium of claim 31, wherein
closing the first channel comprises a whole or partial network
outage that disrupts the first channel.

20

33. The computer-readable medium of claim 31, wherein closing the first channel comprises rebooting a machine associated with the first channel.

5 34. The computer-readable medium of claim 33, wherein the machine comprises a client accessing the file via the first handle.

10 35. The computer-readable medium of claim 33, wherein the machine comprises a server providing access to the file.

36. A computer-executable method for accessing to a file, comprising:

15 obtaining a first handle to a file via a loopback path, the first handle capable of accessing the file in accordance with any rights granted while opening the file;

 obtaining a resume key with which to request a duplicate handle, the duplicate handle capable of accessing the file in any way in which the first handle is capable; and

20 obtaining the duplicate handle to the file and thereafter performing any read or write access to the file via the duplicate handle.

37. The computer-readable medium of claim 36, wherein a client requests the handles and a server provides the handles and wherein the client and the server both reside on a single
5 machine.

38. The computer-readable medium of claim 36, wherein the loopback path comprises a transmission control protocol/internet protocol (TCP/IP) stack.

10

39. The computer-readable medium of claim 36, wherein the resume key comprises a file identifier that identifies an open file, a time stamp, and a process identifier that identifies a process associated with the resume key.

15

40. The computer-readable medium of claim 39, wherein at least part of the resume key provides an index for other state information associated with the file.

20

41. A system for accessing a file, comprising:
a client configured to send a request to open a file,
receive a first handle to the file and a resume key, send a

request for a duplicate handle together with the resume key,
and access the file via the first and duplicate handles; and

a plurality of servers associated with a data store, only
an active one of the plurality of servers controlling the data
5 store at a time, the plurality of servers configured in a
redundant relationship wherein when the active server goes
offline another one of the plurality of servers becomes active
and takes control of the data store, the data store comprising
open files including the file and a data structure including
10 resume keys to obtain duplicate handles for the open files,

wherein after the active server goes offline, the other
server that becomes active is configured to build a data
structure that is included on the other server and to place
resume keys stored in the data store therein, so that the
15 other server may provide a duplicate handle in response to a
request for the duplicate handle from the client.

42. The system of claim 41, wherein the first handle and
the duplicate handle provide identical access to the file.

20

43. A system for accessing a file, comprising:

a client configured to send a request to open a file,
receive a first handle to the file and a resume key, send a
request for a duplicate handle together with the resume key,
5 and access the file via the first and duplicate handles;

a plurality of data stores configured to operate in a
distributed file system; and

a plurality of servers, each server associated with one
of the data stores, each server including a data structure
10 comprising resume keys to use in providing duplicate handles
for the open files included on the data store associated with
the server,

wherein when any of the servers is notified that it will
be taken offline, the server is configured to send data
15 including resume keys from its data structure to another
server and to refer any requests for access to files
associated with the sent data to the other server, the other
server providing access to the file thereafter, the other
server incorporating the data into the data structure included
20 on the other server, so that the other server provides a
duplicate handle in response to a request for the duplicate
handle from the client.

44. The system of claim 43, wherein the first handle and the duplicate handle have the same access rights to the file.